

Recognition of unipolar and generalised split graphs

Colin McDiarmid
cmcd@stats.ox.ac.uk

Nikola Yolov
niklov@cs.ox.ac.uk

April 5, 2016

Abstract

A graph is unipolar if it can be partitioned into a clique and a disjoint union of cliques, and a graph is a generalised split graph if it or its complement is unipolar. A unipolar partition of a graph can be used to find efficiently the clique number, the stability number, the chromatic number, and to solve other problems that are hard for general graphs. We present the first $O(n^2)$ time algorithm for recognition of n -vertex unipolar and generalised split graphs, improving on previous $O(n^3)$ time algorithms.

1 Introduction

1.1 Definition and motivation

A graph is *unipolar* if for some $k \geq 0$ its vertices admit a partition into $k + 1$ cliques $\{C_i\}_{i=0}^k$ so that there are no edges between C_i and C_j for $1 \leq i < j \leq k$. A graph G is a *generalised split graph* if either G or its complement \bar{G} is unipolar. All generalised split graphs are perfect; and Prömel and Steger [PS92] show that almost all perfect graphs are generalised split graphs. Perfect graphs can be recognised in polynomial time [CLV03] and [CCL⁺05], and there are many NP-hard problems which are solvable in polynomial time for perfect graphs, including the stable set problem, the clique problem, the colouring problem, the clique covering problem and their weighted versions [GLS84]. If the input graph is restricted to be a generalised split graph, then there are much more efficient algorithms for the problems above [EW14]. In this paper we address the problem of efficiently recognising generalised split graphs, and finding a witnessing partition.

Previous recognition algorithms for unipolar graphs include [TC85] which achieves $O(n^3)$ running time, [CH12] with $O(n^2m)$ time and [EW14] with $O(nm + nm')$ time, where n and m are respectively the number of vertices and edges of the input graph, and m' is the number of edges added after a triangulation of the input graph. Note that almost all unipolar graphs and almost all generalised split graphs have $(1 + o(1))n^2/4$ edges [MY16]. Further, by testing whether G

or \overline{G} is unipolar, each of the mentioned algorithms above recognises generalised split graphs in $O(n^3)$ time. The algorithm in this paper has running time $O(n^2)$.

This leads to polynomial-time algorithms for the problems mentioned above (stable set, clique, colouring and so on) which have $O(n^{2.5})$ expected running time for a random perfect graph R_n and an exponentially small probability of exceeding this time bound. Here we assume that R_n is sampled uniformly from the perfect graphs on vertex set $[n] = \{1, 2, \dots, n\}$.

1.2 Notation

We use $V(G)$, $E(G)$, $v(G)$ and $e(G)$ to denote V , E , $|V|$ and $|E|$ for a graph $G = (V, E)$. We let $N(v)$ denote the neighbourhood of a vertex v , and let $N^+(v)$ denote $N(v) \cup \{v\}$, also called the closed neighbourhood of v . If $G = (V, E)$ and $S \subseteq V$, then $G[S]$ denotes the subgraph induced by S . Let \mathcal{GS}^+ be the set of all unipolar graphs and let \mathcal{GS} be the set of all generalised split graphs. Fix $G = (V, E) \in \mathcal{GS}^+$. If $V_0, V_1 \subseteq V$ with $V_0 \cap V_1 = \emptyset$ and $V_0 \cup V_1 = V$ are such that V_0 is a clique and V_1 is a disjoint union of cliques, then the ordered pair (V_0, V_1) will be called a *unipolar representation* of G or just a *representation* of G . For each unipolar representation $R = (V_0, V_1)$ we call V_0 the *central clique* of R , and we call the maximal cliques of V_1 the *side cliques* of R . A graph is unipolar iff it has a unipolar representation.

Definition 1.1. Let $R = (V_0, V_1)$ be a unipolar representation of a graph G . A partition \mathcal{B} of $V(G)$ is a *block decomposition* of G with respect to R if the intersection of each part of \mathcal{B} with V_1 is either a side clique or \emptyset .

1.3 Plan of the paper

Assume that G is an input graph throughout. The algorithm for recognising unipolar graphs has three stages. In the first stage a sufficiently large maximal independent set is found. The second stage constructs a partition \mathcal{B} of $V(G)$, such that \mathcal{B} is a block decomposition for some unipolar representation if $G \in \mathcal{GS}^+$. The third stage generates a 2-CNF formula which is satisfiable iff \mathcal{B} is a block decomposition for some unipolar representation. The formula is constructed in such a way that a satisfying assignment of the variables corresponds to a representation of G , and the algorithm returns either a representation of G , or reports $G \notin \mathcal{GS}^+$.

We describe the third stage first (in §2) as it is short and includes a natural transformation to 2-SAT. In §3 we discuss the first stage of finding large independent sets. In §4 we present the second stage when we seek to build a block decomposition. Finally, in §5, we briefly discuss random perfect graphs and algorithms for them using the algorithm described above.

1.4 Data Structures

The most commonly used data type for this algorithm is the set. We assume that the operation $A \cap B$ takes $O(\min(|A|, |B|))$ time, $A \cup B$ takes $O(|A| + |B|)$

time, $A \setminus B$ takes $O(|A|)$ time and $a \in A$ takes $O(1)$ time. These properties can be achieved by using hashtables to implement sets.

Functions will always be of the form $f : [m] \rightarrow A$ for some m , where $[m] = \{1, 2, \dots, m\}$. Therefore functions can be implemented with simple arrays, hence the lookup and assignment operations are assumed to require $O(1)$ time.

2 Verification of Block Decomposition

2.1 2-SAT

Let x_1, \dots, x_n be n boolean variables. A *2-clause* is an expression of the form $y_1 \vee y_2$, where each y_j is a variable, x_i , or the negation of a variable, $\neg x_i$. There are $4n^2$ possible 2-clauses. The problem of deciding whether or not a formula of the form $\psi = \exists x_1 \exists x_2 \dots \exists x_n (c_1 \wedge c_2 \wedge \dots \wedge c_m)$, where each c_j is a 2-clause, is satisfiable is called 2-SAT. The problem 2-SAT is solvable in $O(n + m)$ time – [EIS76] and [APT79], where n is the number of variables and m is the number of clauses in the input formula.

2.2 Transformation to 2-SAT

Let $G = (V, E)$ be a graph with vertex set $V = [n]$. In this subsection we show how to test if a partition of V is a block decomposition for some unipolar representation, in which case we must have $G \in \mathcal{GS}^+$. Let \mathcal{B} be the partition of V we want to test. From each block of \mathcal{B} , we seek to pick out some vertices to form the central clique V_0 of a representation, with the remaining vertices in the blocks forming the side cliques. Suppose that $|\mathcal{B}| = m$, and \mathcal{B} is represented by a surjective function $f : V \rightarrow [m]$, so that $\mathcal{B} = \{f^{-1}[i] : i \in [m]\}$. Let $\{x_v : v \in V\}$ be Boolean variables. We use the procedure **verify** to construct a formula $\psi(x_1, \dots, x_n)$, so that each satisfying assignment of $\{x_v\}$ corresponds to a representation of G .

Procedure verify(G, f):

```

 $\psi := \exists x_1 \exists x_2 \dots \exists x_n$ 
for  $\{u, v\} \in V^{(2)}$  do
  if  $uv \in E$  then
    if  $f(u) = f(v)$  then
      do nothing
    else
       $\psi := \psi \wedge (x_u \vee x_v)$ 
    end if
  else
    if  $f(u) = f(v)$  then
       $\psi := \psi \wedge (x_u \vee x_v) \wedge (\neg x_u \vee \neg x_v)$ 
    else
       $\psi := \psi \wedge (\neg x_u \vee \neg x_v)$ 
    end if

```

end if
end for
return 2-SAT(ψ)

There is an exception: the first time a clause is added to ψ it should be added without the preceding sign for conjunction. The following lemma is easy to check.

Lemma 2.1. *The formula ψ is satisfiable iff \mathcal{B} is a block decomposition for some representation. Indeed an assignment $\Phi : \{x_v : v \in V\} \rightarrow \{0, 1\}$ satisfies ψ if and only if $R = (V_0, V_1)$ is a representation of G and \mathcal{B} is a block decomposition of G with respect to R , where $V_i = \{v \in V : \Phi(v) = 1 - i\}$.*

Proof. Suppose Φ is a satisfying assignment and let V_0, V_1 be as above. If u and v are both in V_0 , then $uv \in E$, since otherwise Φ contains a clause $\neg x_u \vee \neg x_v$. If u and v are in V_1 , then either $uv \in E$ and $f(u) = f(v)$ or $uv \notin E$ and $f(u) \neq f(v)$, because in the other two cases ϕ contains the clause $x_u \vee x_v$. This means that the vertices in V_1 are grouped into cliques by their value of f . For the other direction, it is sufficient to verify that each generated clause is satisfied, which is a routine check. \square

At most a constant number of operations are performed per pair $\{u, v\}$, so $O(n^2)$ time is spent preparing ψ . The formula ψ can have at most 2 clauses per pair $\{u, v\}$, so the length of ψ is also $O(n^2)$, and since 2-SAT can be solved in linear time, the total time for this step is $O(n^2)$.

3 Independent Set

3.1 Maximum Independent Set of a Unipolar Graph


Let $\alpha(G)$ be the maximum size of an independent set in a graph G . Let $G \in \mathcal{GS}^+$ and let R be a unipolar representation of G . Observe that for any representation R of G , the number $s(R, G)$ of side cliques satisfies $s(R, G) \leq \alpha(G) \leq s(R, G) + 1$. We deduce that for every two representations R_1 and R_2 of G we have $|s(R_1, G) - s(R_2, G)| \leq 1$.

If for example G is K_n or its complement, then the number $s(R, G)$ depends on R . However, this is not necessarily the case for all graphs, see Figure 1. It can be shown that the number of n -vertex unipolar graphs with a unique representation is $(1 - e^{-\Theta(n)})|\mathcal{GS}_n^+|$, and that the number of n -vertex unipolar graphs G with a unique representation R and such that $s(G, R) = \alpha(G)$ is $(1 - O(e^{-n^\delta}))|\mathcal{GS}_n^+|$ for a constant $\delta > 0$ [MY16].

3.2 Independent Set Algorithm

It is well known that calculating $\alpha(G)$ for a general graph is NP-hard. For $G \in \mathcal{GS}^+$ let $s(G) = \max_R s(R, G)$, where the maximum is over all representations R of G . For $G \notin \mathcal{GS}^+$ set $s(G) = 0$. In this section we see how to find a maximal independent set I , such that if $G \in \mathcal{GS}^+$, then $|I| \geq s(G) (\geq \alpha(G) - 1)$.

Figure 1: For the graph G_1 on the left, $s(R, G_1) = \alpha(G_1)$ for all representations R , and the graph G_2 on the right, $s(R, G_2) = \alpha(G_1) + 1$ for all representations R .



The idea is to start with G and with $I = \emptyset$; and as long as the remaining graph has two non-adjacent vertices, say v_1 and v_2 , pick $r = 1$ or 2 of these vertices to add to I , and delete from G the closed neighbourhood of the added vertices. We do this in such a way that a given representation R of G yields a representation with r less side cliques, or (only when $r = 2$) with one less side clique and the central clique removed.

Procedure $\text{indep}(G)$:

```

 $I := \emptyset$ ,  $U := V(G)$ 
while  $U \neq \emptyset$  do
  if  $G[U]$  is complete then
    pick an arbitrary  $u \in U$ 
    return  $I \cup \{u\}$ .
  else
    pick  $u_1, u_2 \in U$ , so that  $u_1 u_2 \notin E(G)$ .
     $U_1 := (N^+(u_1) \setminus N^+(u_2)) \cap U$ 
     $U_2 := (N^+(u_2) \setminus N^+(u_1)) \cap U$ 
    if  $G[U_1]$  is complete then
      if  $G[U_2]$  is complete then
         $I := I \cup \{u_1, u_2\}$ 
         $U := U \setminus (N^+(u_1) \cup N^+(u_2))$ 
      else
         $I := I \cup \{u_1\}$ 
         $U := U \setminus N^+(u_1)$ 
      end if
    else
      if  $G[U_2]$  is complete then
         $I := I \cup \{u_2\}$ 
         $U := U \setminus N^+(u_2)$ 
      else
         $I := I \cup \{u_1, u_2\}$ 
         $U := U \setminus (N^+(u_1) \cup N^+(u_2))$ 
      end if
    end if
  end if
end while

```

return I

Observe that the main body of the $\text{indep}(G)$ procedure is a **while** loop. An alternative way of seeing the algorithm is that instead of the loop there is a recursive call to $\text{indep}(G[U])$ at the end of the iteration and the procedure returns the union of the vertices found during this iteration and the recursively retrieved set. A recursive interpretation is clearer to work with for inductive proofs.

3.3 Correctness

Lemma 3.1. *Procedure $\text{indep}(G)$ always returns a maximal independent set I .*

Proof. This is easy to see, since each vertex deleted from U is adjacent to a vertex put in I . \square

Lemma 3.2. *If $\text{indep}(G)$ returns I , then $|I| \geq s(G)$.*

Proof. If $G \notin \mathcal{GS}^+$, then the statement holds, because $s(G) = 0$. From now on assume that $G \in \mathcal{GS}^+$. We argue by induction on $v(G)$. It is trivial to see that the lemma holds for $v(G) = 1$. Let $v(G) > 1$ and assume that the lemma holds for smaller graphs. If G is complete, then $|I| = 1 = s(G)$.

Fix an arbitrary unipolar representation R of G . We show that $|I| \geq s(R, G)$. If G is not complete, then the procedure selects two non-adjacent vertices u and v . The vertices u and v are either in different side cliques or one of them is in the central clique and the other is in a side clique.

We start with the case when u and v are contained in side cliques. After inspecting their neighbourhoods, the algorithm removes from U either one or both of them along with their neighbourhood. Suppose that it removes r of them, where r is 1 or 2. Let G' and R' be the graph and the representation induced by the remaining vertices. By the induction hypothesis, if I' is the recursively retrieved set, then $|I'| \geq s(R', G')$. Let I be the independent set returned at the end of the algorithm, so that $|I'| + r = |I|$. Both u and v see all the vertices in their corresponding side clique, and see no vertices from different side cliques, so after removing r of them with their neighbours, the number of side cliques in the representation decreases by precisely r , and hence $s(R, G) = s(R', G') + r$. Now $|I| = |I'| + r \geq s(R', G') + r = s(R, G)$.

Now w.l.o.g. assume that u belongs to a side clique and v belongs to the central clique. Then $N^+(u)$ contains the side clique of u and perhaps parts of the central clique. Therefore, $N^+(u) \setminus N^+(v)$ is a subset of the side clique of u , and hence it is a clique. If $N^+(v) \setminus N^+(u)$ is a not clique, then the algorithm continues recursively with $G[V \setminus N^+(u)]$; and using the same arguments as above with $r = 1$, we guarantee correct behaviour. Now assume that $N^+(v) \setminus N^+(u)$ is a clique. Then $N^+(v) \setminus N^+(u)$ can intersect at most one side clique, because the vertices in different side cliques are not adjacent. In this case $N^+(v) \cup N^+(u)$ completely covers the side clique of u , completely covers the central clique, and it may intersect one additional side clique. Hence $s(R, G) = s(R', G') + 1$ or $s(R, G) = s(R', G') + 2$, where G' and R' are the induced graph and

representation after the removal of $N^+(v) \cup N^+(u)$. If I' is the recursively obtained independent set, from the induction hypothesis we deduce that $|I| = |I'| + 2 \geq s(R', G') + 2 \geq s(R, G)$. \square

3.4 Time Complexity

In this form the algorithm takes more than $O(n^2)$ time, because checking whether an induced subgraph is complete is slow. However, we can maintain a set of vertices, C , which we have seen to induce a complete graph. We will create an efficient procedure to check if a subgraph is complete, and to return some additional information to be used for future calls if the subgraph is not complete.

Procedure `antiedge`(G, U, C):

```

 $C' := C$ 
for  $v \in U \setminus C$  do
  if  $C' \setminus N^+(v) = \emptyset$  then
     $C' := C' \cup \{v\}$ 
  else
    let  $u \in C' \setminus N^+(v)$ 
    return  $(uv, C')$ 
  end if
end for
return  $(False, C')$ .

```

The following lemma summarises the behaviour of `antiedge`.

Lemma 3.3. *Let $C \subseteq U$ and suppose that $G[C]$ is complete. If $G[U]$ is complete, then `antiedge`(G, U, C) returns $(False, U)$; if not, then it returns (uv, C') such that*

1. $uv \in C' \times (U \setminus C') - E(G)$ i.e. $u \in C', v \in U \setminus C', uv \notin E(G)$
2. $C \subseteq C' \subseteq U$
3. $G[C']$ is complete

Proof. Easy checking. \square

Procedure `indep`(G):

```

 $I := \emptyset; U := V(G)$ 
 $e := False; C := \emptyset$ 
 $(e, C) := \text{antiedge}(G, U, C)$ 
while  $U \neq \emptyset$  do
  if  $e = False$  then
    pick an arbitrary  $u \in U$ 
    return  $I \cup \{u\}$ .
  else
    Assume that  $e = u_1u_2; u_1, u_2 \in U$ 

```

```

 $U_1 := (N^+(u_1) \setminus N^+(u_2)) \cap U$ 
 $U_2 := (N^+(u_2) \setminus N^+(u_1)) \cap U$ 
 $(e_1, C_1) := \text{antiedge}(G, U_1, C \cap U_1)$ 
 $(e_2, C_2) := \text{antiedge}(G, U_2, C \cap U_2)$ 
if  $e_1 = \text{False}$  then
  if  $e_2 = \text{False}$  then
     $I := I \cup \{u_1, u_2\}$ 
     $U := U \setminus (N^+(u_1) \cup N^+(u_2))$ 
     $(e, C) := \text{antiedge}(G, U, \emptyset)$ 
  else
     $I := I \cup \{u_1\}$ 
     $U := U \setminus N^+(u_1)$ 
     $C := C_2; e := e_2$ 
  end if
else
  if  $e_2 = \text{False}$  then
     $I := I \cup \{u_2\}$ 
     $U := U \setminus N^+(u_2)$ 
     $C := C_1; e := e_1$ 
  else
     $I := I \cup \{u_1, u_2\}$ 
     $U := U \setminus (N^+(u_1) \cup N^+(u_2))$ 
     $(e, C) := \text{antiedge}(G, U, \emptyset)$ 
  end if
end if
end while
return  $I$ 

```

Lemma 3.4. *Let U and C be the sets stored in the respective variables at the beginning of an iteration of the main loop of the modified **indep**, and let U' and C' be the sets stored at the beginning of the next iteration, if the algorithm does not terminate meanwhile. The following loop invariants hold:*

(I1) $G[C]$ is complete and $C \subseteq U$,

(I2) $U' \setminus C' \subseteq U \setminus C$.

A loop invariant is a condition which is true at the beginning of each iteration of a loop.

Proof. Observe that the initial values of U and C , which are V and \emptyset respectively, guarantee by Lemma 3.3 that the values after the call to **antiedge** satisfy condition (I1). Therefore (I1) holds for the first iteration. Concerning future iterations, observe that (I1) guarantees the precondition of Lemma 3.3, which it turn guarantees (I1) for the next iteration. We deduce that (I1) does indeed give a loop invariant. By proving this we have proved that the preconditions of Lemma 3.3 are always met; and so we can use Lemma 3.3 throughout.

If $e = \text{False}$, then there is no next iteration, hence condition (I2) is automatically correct. Now assume that $e = u_1 u_2$. Depending on e_1 and e_2 there are two cases for how many vertices are excluded. Case 1: one vertex is excluded. W.l.o.g. assume that u_1 is excluded, so $U' = U \setminus N^+(u_1)$, $C \cap U_2 \subseteq C'$ and $C \subseteq N^+(u_1) \cup N^+(u_2)$. Then

$$\begin{aligned} U' \setminus C' &\subseteq U' \setminus (C \cap U_2) \\ &= (U \setminus N^+(u_1)) \setminus (C \cap (N^+(u_2) \setminus N^+(u_1)) \cap U) \\ &= U \setminus [N^+(u_1) \cup (C \cap (N^+(u_2) \setminus N^+(u_1)))] \\ &= U \setminus [N^+(u_1) \cup C] \subseteq U \setminus C. \end{aligned}$$

Case 2: two vertices are excluded. Now $U' = U \setminus (N^+(u_1) \cup N^+(u_2))$, and

$$U' \setminus C' \subseteq U' \subseteq U \setminus C.$$

We have shown that condition (I2) holds at the start of the next iteration, and so it gives a loop invariant as claimed. \square

A vertex v is *absorbed* if it is processed during the loop of **antiedge** and then appended to the result set, C' .

Corollary 3.5. *A vertex can be absorbed once at most.*

Proof. Let U, C, U' and C' be as before. Observe that if, during the iteration, vertex v is absorbed in a call to **antiedge**, then $v \in U \setminus C$ and $v \notin U' \setminus C'$. The Corollary now follows from the second invariant in Lemma 3.4. \square

Lemma 3.6. *The procedure **indep**(G) using **antiedge** takes $O(n^2)$ time.*

Proof. The total running time of each iteration of the main loop of **indep**(G) besides calling **antiedge** is $O(n)$. The set U decreases by at least one vertex on each iteration, so the time spent outside of **antiedge** is $O(n^2)$.

From Corollary 3.5 at most n vertices are absorbed and $O(n)$ steps are performed each time, so in total $O(n^2)$ time is spent in all calls to **antiedge** for absorbing vertices.

Assume that v is processed in **antiedge** for the first time, but it is not absorbed and it is tested against a set C_1 . Since v is not absorbed, we may assume that **antiedge** has returned the pair of vertices vu . At least one of u and v is removed (along with its neighbourhood) from U and moved to I . If v is removed from U , then no more time can be spent on it in **antiedge**, hence the total time spent on v in **antiedge** is $O(n)$. Now assume that u is removed. We have that $C_1 \subseteq N^+(u)$ and each vertex in $N^+(u)$ is removed from U . Hence, if v is processed again in **antiedge**, it will be tested against a set C_2 with $C_1 \cap C_2 = \emptyset$, and therefore $|C_1| + |C_2| = |C_1 \cup C_2| = O(n)$. As we saw before, if v is absorbed or removed from U , then it cannot be processed again in **antiedge**; and thus the running time spent on v is again $O(n)$. If v is not removed from U , then C_2 is removed from U . Hence, if v is processed again in **antiedge**, v

will be tested against a set C_3 with $|C_1| + |C_2| + |C_3| = |C_1 \cup C_2 \cup C_3| = O(n)$, and so on. Thus, we see that over all these tests, each vertex is tested at most once for adjacency to v , and so the total time spent on v is $O(n)$. \square

4 Building Blocks and Recognition

4.1 Block Creation Algorithm

In this subsection we present a short algorithm for creating a partition of $V(G)$ using an independent set I and then checking if this partition is a block decomposition using the procedure `verify` from Section 2.

Procedure `test`(G, I):

```

 $U := V(G); t := 0; f := \emptyset$ 
for  $i \in I$  do
  for  $v \in N^+(i) \cap U$  do
     $f(v) := t$ 
  end for
   $U := U \setminus N^+(i)$ 
   $t := t + 1$ 
end for
for  $v \in U$  do
   $f(v) := t$ 
end for
return verify( $G, f$ )

```

Lemma 4.1. *Suppose that $I \subseteq V(G)$ is an independent set with $|I| \geq s(G) - 1$ and $V_0 \cap I = \emptyset$ for some unipolar representation $R = (V_0, V_1)$ of G . Then `test`(G, I) returns *True*.*

Proof. On each step of the main loop a vertex from $i \in I$ is selected. Since $V_0 \cap I = \emptyset$, the vertex i is a part of some side clique, say C . Now $C \cap N^+(j) = \emptyset$ for each $j \in I \setminus \{i\}$, so $C \subseteq U$. Also $C \subseteq N^+(i)$, and hence $C \subseteq N^+(i) \cap U$. Vertex i does not see vertices from other side cliques, so $N^+(i) \cap U$ is correctly marked as a separate block.

Since $|I| \geq s(G) - 1$, at most one side clique is not represented in I . If there is an unrepresented side clique, say C , then none of the previously created blocks can claim any vertex from it, and hence $C \subseteq U$. We have shown that when the main loop ends, either $U \cap V_1 = \emptyset$ or $U \cap V_1$ is a side clique; so U is correctly marked as a separate block. The set U also contains all remaining vertices, so f is partition of V into blocks, and hence `verify`(G, f) will return *True*. \square

4.2 Block Decomposition Algorithm

By Lemma and 3.1 and 3.2, `indep`(G) returns a maximal independent set I of size at least $s(G)$. Thus, Lemma 4.1 suggests a naive algorithm for recognition

for \mathcal{GS}^+ – try $\text{test}(G, I \setminus i)$ for each $i \in I$ and return *True* if any attempt succeeds. The proposed algorithm is correct, since $|I \cap V_0| \leq 1$. The running time is $O(|I|n^2) = O(n^3)$, while we aim for $O(n^2)$. However, with relatively little effort we can localise $I \cap V_0$ to at most 2 candidates from I .

Procedure `blocks`(G, I):

```

 $C := I$ 
for  $v \in V(G)$  do
  if  $|N^+(v) \cap I| = 2$  then
     $C := C \cap N^+(v)$ 
  end if
end for
if  $|C| = 1$  then
  return  $\text{test}(G, I \setminus C)$ 
else if  $|C| = 2$  then
  Assume that  $C = \{c_1, c_2\}$ 
  return  $\text{test}(G, I \setminus \{c_1\}) \vee \text{test}(G, I \setminus \{c_2\})$ 
else
  return  $\text{test}(G, I)$ 
end if

```

Procedure `recognise`(G):

```

return blocks( $G, \text{indep}(G)$ )

```

4.3 Correctness

Lemma 4.2. *The procedure `recognise`(G) returns *True* iff $G \in \mathcal{GS}^+$.*

Proof. First assume that $G \in \mathcal{GS}^+$ and let $R = (V_0, V_1)$ be an arbitrary representation of G . Let $I = \text{indep}(G)$. By Lemma 3.2, $|I| \geq s(G) \geq s(R, G)$. Since V_0 is a clique and I is an independent set, we have $|V_0 \cap I| \leq 1$.

Case 1: $V_0 \cap I = \emptyset$. Observe that `blocks` returns $\text{test}(G, I')$, where I' is either I or $I \setminus \{v\}$ for some $v \in I$, hence $|I'| \geq |I| - 1 \geq s(G) - 1$; $I' \subseteq I \subseteq V_1$, so $\text{test}(G, I') = \text{True}$ from Lemma 4.1.

Case 2: $V_0 \cap I = \{c\}$. `Blocks` starts by calculating the set C , where $C = I$ if there is no $v \in V$ with $|N^+(v) \cap I| = 2$, and otherwise

$$C = \bigcap \{N^+(v) \cap I : v \in V(G), |N^+(v) \cap I| = 2\}.$$

Assume that $|N^+(v) \cap I| = 2$ for some $v \in V$. If $v \in V_0$, then $c \in N^+(v)$, because V_0 is a clique. If $v \in V_1$, then $N^+(v)$ can intersect at most one vertex from $I \cap V_1$ and at most one vertex from $I \cap V_0 = \{c\}$ and since $|N^+(v) \cap I| = 2$, we have $c \in N^+(v)$. For each $v \in V$ if $|N^+(v) \cap I| = 2$, then $c \in N^+(v) \cap I$, so c belongs to their intersection. If no $v \in V$ exists with $|N^+(v) \cap I| = 2$, then $C = I$, but $c \in I$, so again $c \in C$. We deduce that if $V_0 \cap I = \{c\}$, then $c \in C$ and $|C| > 0$.

If $|C| = 1$ or $|C| = 2$ then $\text{test}(G, I \setminus \{i\})$ is tested individually for each vertex $i \in C$, but $c \in C$ and $\text{test}(G, I \setminus \{c\}) = \text{True}$ by Lemma 4.1.

If $|C| > 2$, then there is no $v \in V$ with $|N^+(v) \cap I| = 2$. Either $|I| = s(R, G)$ or $|I| = s(R, G) + 1$, so either all side cliques are represented by vertices of I , or at most one is not represented, say S . We can handle both cases simultaneously by saying that $S = \emptyset$ in the former case. We have that I is a maximal independent set, but no vertex of $I \setminus \{c\}$ can see a vertex of S , because they belong to different side cliques, so c is connected to all vertices of S and therefore $\{c\} \cup S$ is a clique. Let $T = N(c) \cap (V_1 \setminus S)$. Then $|N^+(v) \cap I| = 2$ for each $v \in T$, but no such vertex exists by assumption, so $T = \emptyset$. Now $N(c) \cap V_1 = S$, and V_1 is a union of disjoint cliques, so $V_1 \cup \{c\}$ is also a union of disjoint cliques. Hence $R' = (V_0 \setminus \{c\}, V_1 \cup \{c\})$ is a representation of G , so from Lemma 4.1 $\text{test}(G, I) = \text{True}$.

On the contrary, if $G \notin \mathcal{GS}^+$, then there is no representation for G , hence test cannot generate a block decomposition of G , and therefore test will return *False*. \square

4.4 Time Complexity

Lemma 4.3. *$\text{recognise}(G)$ takes $O(n^2)$ time.*

Proof. The procedure test loops over a subset of V and intersects two subsets of V , so the time for each step is bounded by $O(n)$, and since the number of steps is $O(n)$, $O(n^2)$ time is spent in the loop. Then it performs one more operation in $O(n)$ time, so the total time spent for preparation is $O(n^2)$. Then test calls verify , which takes $O(n^2)$ time, so the total running time of test is $O(n^2)$.

While building C , blocks handles $O(n)$ sets with size $O(n)$, so it spends $O(n^2)$ time in the first stage. Depending on the size of C , blocks calls test once or twice, but in both cases it takes $O(n^2)$ time, so the total running time of blocks is $O(n^2)$. The total time spent for recognition is the time spent for blocks plus the time spent for indep, and since both are $O(n^2)$, the total running time for recognition is $O(n^2)$. \square

5 Algorithms for random perfect graphs

Grötschel, Lovász, and Schrijver [GLS84] show that the stable set problem, the clique problem, the colouring problem, the clique covering problem and their weighted versions are computable in polynomial time for perfect graphs. The algorithms rely on the Lovász sandwich theorem, which states that for every graph G we have $\omega(G) \leq \vartheta(\overline{G}) \leq \chi(G)$, where $\vartheta(G)$ is the Lovász number. The Lovász number can be approximated via the ellipsoid method in polynomial time, and for perfect graphs we know that $\omega(G) = \chi(G)$, hence $\vartheta(G)$ is an integer and its precise value can be found. Therefore $\chi(G)$ and $\omega(G)$ can be found in polynomial time for perfect graphs, though these are NP-hard problems for general graphs. Further, $\alpha(G)$ and $\overline{\chi}(G)$ (the clique covering number) can be

computed from the complement of G (which is perfect). The weighted versions of these parameters can be found in a similar way using the weighted version of the Lovász number, $\vartheta_w(G)$.

These results tell us more about computational complexity than algorithm design in practice. On the other hand, the problems above are much more easily solvable for generalised split graphs. We know that the vast majority of the n -vertex perfect graphs are generalised split graphs [PS92]. One can first test if the input perfect graph is a generalised split graph using the algorithm in this paper and if so, apply a more efficient solution.

Eschen and Wang [EW14] show that, given a generalised split graph G with n vertices together with a unipolar representation of G or \overline{G} , we can efficiently solve each of the following four problems: find a maximum clique, find a maximum independent set, find a minimum colouring, and find a minimum clique cover.

It is sufficient to show that this is the case when G is unipolar, as otherwise we can solve the complementary problem in the complement of G . Finding a maximum size stable set and minimum clique cover in a unipolar graph is equivalent to determining whether there exists a vertex in the central clique such that no side clique is contained in its neighbourhood, which is trivial and can be done very efficiently. Suppose there are k side cliques. If there is such a vertex v , then a maximum size stable set (of size $k + 1$) consists of v and from each side clique a vertex not adjacent to v , and a minimum size clique cover is formed by the central clique and the k side cliques. If not, then a maximum size stable set (of size k) consists of a vertex from each side clique, and a minimum clique cover is formed by extending the k side cliques to maximal cliques (which then cover C_0).

Let us focus on finding a maximum clique and minimum colouring of a unipolar graph G with a representation R . If R contains k side cliques, C_1, \dots, C_k , then

$$\omega(G) = \chi(G) = \max\{\omega(G[C_0 \cup C_i])\}_{i=1}^k = \max\{\chi(G[C_0 \cup C_i])\}_{i=1}^k,$$

where C_0 is the central clique. Therefore, in order to find a maximum clique or a minimum colouring, it is sufficient to solve the corresponding problem in each of the co-bipartite graphs induced by the central clique and a side clique. The vertices outside a clique in a co-bipartite graph form a cover in the complementary bipartite graph, and the vertices coloured with the same colour in a proper colouring of a co-bipartite graph form a matching in the complementary bipartite graph. By König's theorem it is easy to find a minimum cover using a given maximum matching, and therefore finding a maximum clique and a minimum colouring in a co-bipartite graph is equivalent to finding a maximum matching in the complementary bipartite graph. For colourings, we explicitly find a minimum colouring in each co-bipartite graph $G[C_0 \cup C_i]$, and such colourings can be fitted together using no more colours, since C_0 is a clique cutset. Assume that $\overline{G[C_0 \cup C_i]}$ contains n_i vertices and m_i edges, so each $n_i \leq n$ and $\sum_i m_i \leq |C_0|(n - |C_0|) \leq n^2/4$. We could use the Hopcroft–Karp algo-

rithm for maximum matching in $O((|E| + |V|)\sqrt{|V|})$ time to find time bound $\sum_i O((m_i + n_i)\sqrt{n_i}) = O((n + m)\sqrt{n}) = O(n^{2.5})$.

The approach of Eschen and Wang [EW14] is very similar, and they give more details, but unfortunately there is a mistake with their analysis, and a corrected version of their analysis yields $O(n^{3.5}/\log n)$ time, instead of the claimed $O(n^{2.5}/\log n)$. In order to see the mistake consider the case when the input graph is a split graph with an equitable partition.

Given a random perfect graph R_n , we run our recognition algorithm in time $O(n^2)$. If we have a generalised split graph, with a representation, we solve each of our four optimisation problems in time $O(n^{2.5})$, if not, which happens with probability $e^{-\Omega(n)}$, we run the methods from [GLS84]. This simple idea yields a polynomial-time algorithm for each problem with low expected running time, and indeed the probability that the time bound is exceeded is exponentially small.

References

- [APT79] Bengt Aspvall, Michael F. Plass, and Robert Endre Tarjan. A linear-time algorithm for testing the truth of certain quantified boolean formulas. *Inf. Process. Lett.*, 8(3):121–123, 1979.
- [CCL⁺05] Maria Chudnovsky, Gérard Cornuéjols, Xinming Liu, Paul Seymour, and Kristina Vušković. Recognizing Berge graphs. *Combinatorica*, 25(2):143–186, 2005.
- [CH12] Ross Churchley and Jing Huang. Solving partition problems with colour-bipartitions. *Graphs and Combinatorics*, pages 1–12, 2012.
- [CLV03] Gérard Cornuéjols, Xinming Liu, and Kristina Vušković. A polynomial algorithm for recognizing perfect graphs. In *Foundations of Computer Science, 2003. Proceedings. 44th Annual IEEE Symposium on*, pages 20–27. IEEE, 2003.
- [EIS76] S. Even, A. Itai, and A. Shamir. On the complexity of timetable and multicommodity flow problems. *SIAM Journal on Computing*, 5(4):691–703, 1976.
- [EW14] Elaine M. Eschen and Xiaoqiang Wang. Algorithms for unipolar and generalized split graphs. *Discrete Appl. Math.*, 162:195–201, January 2014.
- [GLS84] Martin Grötschel, Laszlo Lovász, and Alexander Schrijver. Polynomial algorithms for perfect graphs. *North-Holland Mathematics Studies*, 88:325–356, 1984.
- [MY16] Colin McDiarmid and Nikola YOLOV. Random perfect graphs. *In preparation*, 2016+.

- [PS92] Hans Jürgen Prömel and Angelika Steger. Almost all Berge graphs are perfect. *Combinatorics, Probability and Computing*, 1(01):53–79, 1992.
- [TC85] RI Tyshkevich and AA Chernyak. Algorithms for the canonical decomposition of a graph and recognizing polarity. *Izvestia Akad. Nauk BSSR, ser. Fiz.-Mat. Nauk*, 6:16–23, 1985.